

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants	:	Jewett, et al.
Appl. No.	:	09/927,894
Filed	:	August 10, 2001
For	:	ARCHITECTURE FOR PROVIDING BLOCK-LEVEL STORAGE ACCESS OVER A COMPUTER NETWORK
Examiner	:	Sargon N. Nano
Group Art Unit	:	2157

REPLY BRIEF

United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

This Reply Brief is responsive to the Examiner's Answer in the present application, and supplements the arguments made in Appellants' Appeal Brief (as amended). Because the Examiner's Answer does not raise any new issues with respect to many of the claims at issue, this Reply Brief only discusses some of these claims. More specifically, this Reply Brief focuses on the claims for which the Examiner is now relying on a portion of Wang (U.S. Pat. 6,834,326) that was not previously relied upon to reject the claim.

Independent Claim 1

In the Appeal Brief, Appellants argued that Wang does not explicitly or inherently disclose the limitations of Claim 1 italicized below.

1. A block-level shared network storage system, comprising:
a storage server comprising an array of disk drives, and comprising
a processor that runs a device driver to provide block-level access to data
stored on the array of disk drives; and

a host computer coupled to the storage server by at least one computer network;

wherein the host computer and the storage server perform input/output (I/O) operations over the at least one network using multiple, concurrent logical connections, each logical connection being between the host computer and the storage server over the at least one computer network, such that a first I/O operation is executed over a first logical connection while a second I/O operation is executed over a second logical connection.

As a preliminary matter, Appellants wish to point out that the Examiner has not clearly indicated what in Wang he is treating as the “storage server” recited in Claim 1.

In response to Appellants’ arguments, the Examiner now points to Wang’s description of how a RAID controller can multicast read and write commands to a group of disk drives (referred to as a “multicast group”) that are part of a common RAID volume. In connection with this description, the Examiner asserts that Wang “clearly discloses that multiple logical connections are established between the host (controller) **and multiple disk drives** via a multicasting scheme.” Examiner’s Answer at page 13, first full paragraph (emphasis added), citing col. 11, line 50 – col. 12 and Figures 6 and 8a of Wang.

Even if the Examiner’s assertion is accurate (which Appellants do not concede), Wang still does not explicitly or inherently disclose all of the limitations italicized above. In this regard, the disk drives that are part of a RAID volume or multicast group in Wang are not collectively “a storage server,” much less a storage server as defined in the first subparagraph of Claim 1. Thus, even if Wang’s use of multicasting involves concurrently maintaining a logical connection to each of these disk drives (as the Examiner apparently contends), there still would not be “multiple, concurrent logical connections, each logical connection being between the host computer **and the storage server**” as recited in Claim 1. Nothing in Wang suggests otherwise.

For this reason, and the reasons explained in the Appeal Brief, Appellants respectfully submit that the anticipation rejection of Claim 1 is improper.

Dependent Claim 5

Claim 5 depends from Claim 1, and adds the following limitations: “the host computer is programmed to divide an I/O operation into multiple constituent I/O operations, and to perform

the multiple constituent I/O operations in parallel over respective logical connections of said multiple, concurrent logical connections.” The Examiner now points to column 10, line 66 to column 11, line 29 of Wang in connection with this claim. Examiner’s Answer at page 4, last paragraph. Nothing in this newly cited portion of Wang, however, discloses either the division of an I/O operation into multiple constituent I/O operations, or the performance of multiple constituent I/O operations in parallel over respective logical connections. Indeed, the cited portion of Wang does not appear to describe what happens at the level of I/O operations, but rather focuses on higher level operations.

The Examiner also asserts that “Wang discloses a host or a CPU that sends a command using a SCSI bus, which allows the conducting of multiple block logical read/write operation [sic] in parallel.” Examiner’s Answer at page 13, last paragraph. Even if this statement is true, however, it does not follow that Wang’s CPU “is programmed to divide an I/O operation into multiple constituent I/O operations, and to perform the multiple constituent I/O operations in parallel over respective logical connections of said multiple, concurrent logical connections” as claimed.

The Examiner also points to column 4, lines 40-61 and column 6, line 44 of Wang. In connection with this disclosure, the Examiner states the following: “Wang discloses the implementation of RAID stripes over SCSI where multiple commands are sent to multiple devices to read or write data.” Examiner’s Answer at page 13, last paragraph. Even if this statement is true, however, it does not follow that the limitations of Claim 5 are disclosed. In this regard, the sending of multiple commands to multiple devices does not imply the existence of a host computer that is “programmed to divide an I/O operation into multiple constituent I/O operations, and to perform the multiple constituent I/O operations in parallel over respective logical connections of said multiple, concurrent logical connections,” in the context of the limitations of Claim 1.

For these reasons, and the other reasons explained in the Appeal Brief, Appellants respectfully submit that the anticipation rejection of Claim 5 is improper.

Dependent Claim 7

Claim 7 depends from Claim 1, and adds the following: "the storage server is configurable to provide multiple storage partitions, each of which may be allocated to a different host computer." In connection with this claim, the Examiner states the following: "Wang discloses different types of file systems using different hardware, for instance data base vendors implement their own internal data structure (partitions) for various storage devices, therefore Wang discloses the allocation of different storage partitions (see col. 10, line 66 to col. 11, line 29)." Even if this newly cited portion of Wang inherently discloses a storage server that provides multiple partitions as the Examiner contends (which Appellants do not concede), it does not follow that each such partition can be "allocated to a different host computer" as claimed. Nothing in Wang suggests otherwise. Thus, the newly cited portion of Wang does not support the anticipation rejection of Claim 7.

Dependent Claim 18

Claim 18 depends from Claim 1 and adds the following: "the host computer and the storage server each include two network interfaces that provide redundant network connections between the host computer and the storage server." In connection with this claim, the Examiner now points to column 33, lines 24-38 of Wang. Examiner's Answer at page 14, last two paragraphs. The cited portion of Wang mentions the possibility of adding a second or third network interface to "the switch PC" for purposes of communicating with the disk drives. Regardless of whether this augmented switch PC is treated as the "host computer" or the "storage server," however, the limitations of Claim 18 still are not met. This is because the claim requires that the host computer and the storage server "each include two network interfaces," and requires that these network interfaces "provide redundant network connections between the host computer and the storage server." Nothing in Wang suggests that the addition of a second or third network interface to the switch PC would result in such a configuration.

Independent Claim 19

Claim 19 reads as follows:

19. A system for storing data for host computers, comprising:
a plurality of storage servers connected to a network, each storage server comprising an array of disk drives, an array controller, and a processor;
a plurality of host computers connected to the network and programmed to store data on the storage servers; and
at least one switch which interconnects the plurality of storage servers with the plurality of host computers;
wherein each host computer is programmed to open multiple concurrent socket connections over the network to the storage servers and to perform input/output operations in parallel over the multiple concurrent socket connections.

In the Appeal Brief, Appellants argued that Wang does not disclose the limitations italicized above. The Answer does not clearly indicate whether the Examiner is now relying on Wang's description of multicasting (col. 11, lines 50 – col. 12) in connection with this claim language. See Examiner's Answer at page 13, lines 1-13. To the extent the Examiner may now be relying on this description, Appellants respectfully submit that the multicasting description does not support the rejection. In this regard, a multicast group of disk drives is not a "storage server," much less a storage server as defined in Claim 19. Thus, even if Wang's use of multicasting involves concurrently maintaining a socket connection to each of these disk drives (which Appellants do not concede), the italicized limitations of Claim 19 still would not be met.

For this reason, and the other reasons explained in the Appeal Brief, Appellants respectfully submit that the anticipation rejection of Claim 19 is improper.

Dependent Claim 22

To the extent that Claim 22 recites the same limitations as Claim 5, Appellants submit that the rejection of Claim 22 is improper for the additional reasons explained above for Claim 5.

Dependent Claim 24

To the extent that Claim 24 recites the same limitations as Claim 7, Appellants submit that the rejection of Claim 24 is improper for the additional reasons explained above for Claim 7.

Independent Claim 28

Claim 28 reads as follows:

28. A method of performing input/output operations, comprising:
establishing first and second TCP/IP connections between a host computer and a block-level storage server over one or more computer networks;
performing a first input/output operation over the first TCP/IP connection concurrently with performing a second input/output operation over the second TCP/IP connection, each of said input/output operations comprising a transfer of input/output data between the host computer and the storage server; and
maintaining the first and second TCP/IP connections in a persistent state such that each TCP/IP connection may be used to perform additional input/output operations.

In the Appeal Brief, Appellants argued that the rejection is improper because Wang does not explicitly or inherently disclose the following limitations: “performing a first input/output operation over the first TCP/IP connection [between a host computer and a block-level storage server] concurrently with performing a second input/output operation over the second TCP/IP connection [between the host computer and the block-level storage server], each of said input/output operations comprising a transfer of input/output data between the host computer and the storage server.”

The Answer does not clearly indicate whether the Examiner is now relying on Wang’s description of multicasting (col. 11, lines 50 – col. 12) in connection with this claim language. To the extent the Examiner may now be relying on this description (as with Claim 1), Appellants respectfully submit that the description does not support the rejection. In this regard, a multicast group of disk drives is not a “block-level storage server.” Thus, even if Wang’s use of multicasting inherently involves concurrently maintaining a TCP/IP connection to each of these disk drives (which Appellants do not concede), the above-quoted limitations of Claim 28 still are not met. Thus, Wang’s multicasting description does not support the rejection of Claim 28.

Independent Claim 34

Claim 34 reads as follows:

34. A method of executing an input/output (I/O) request received from a user-level process running on a host computer, comprising:
on the host computer, dividing the I/O request into multiple constituent I/O operations; and
performing the multiple constituent I/O operations in parallel over multiple, respective logical network connections between the host computer and a target storage server such that I/O data is transferred between the host computer and the storage server over each of the logical network connections.

In connection with this claim, the Examiner's Answer states the following:

"Wang discloses a host or a CPU that sends a command using a SCSI bus, which allows the conducting of multiple block logical read/write operation [sic] in parallel. Moreover, Wang discloses the implementation of RAID stripes over SCSI where multiple commands are sent to multiple devices to read or write data (see col. 4, lines 40-61 and col. 6, line 44)." See Examiner's Answer beginning at last full paragraph on page 13.

Even if the above characterizations of Wang are accurate, it does not follow that all of the limitations of Claim 34 are explicitly or inherently disclosed. In this regard, neither the execution of multiple block logical read/write operations in parallel, nor the implementation of RAID stripes over SCSI, inherently involves a host computer that divides an I/O request into multiple I/O operations as claimed. The Examiner has not provided any evidence or reasoning to suggest otherwise. See MPEP 2112 ("In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art.") (citation omitted).

Further, even if there are scenarios in which the Wang's RAID controller divides an I/O request into multiple I/O operations that are handled by different disk drives, the limitations of the second subparagraph of the claim still would not be met. For instance, in such a scenario, the constituent I/O operations would not be performed "over multiple, respective logical network connections between the host computer and a **target storage server**" as claimed.

Appl. No. : 09/927,894
Filed : August 10, 2001

For these reasons, and the reasons explained in the Appeal Brief, Appellants respectfully submit that the anticipation rejection of Claim 34 is improper.

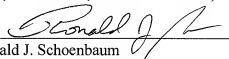
Conclusion

In view of the foregoing, Appellants respectfully submit that the anticipation rejections of Claims 1-37 and 54-59 are improper, and request that these rejections be reversed.

Respectfully submitted,

KNOBBE, MARTENS, OLSON & BEAR, LLP

Dated: 1-12-07

By: 
Ronald J. Schoenbaum
Reg. No. 38,297
Knobbe, Martens, Olson & Bear, LLP
2040 Main Street, 14th Floor
Irvine, CA 92614
949-721-2950
Customer No. 20,995